

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ **BLACK BORDERS**
- ☐ **IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- ☐ **FADED TEXT OR DRAWING**
- ☐ **BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- ☐ **SKEWED/SLANTED IMAGES**
- ☐ **COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- ☐ **GRAY SCALE DOCUMENTS**
- ☐ **LINES OR MARKS ON ORIGINAL DOCUMENT**
- ☐ **REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- ☐ **OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

(12) UK Patent Application (19) GB (11) 2 337 905 (13) A

(43) Date of A Publication 01.12.1999

(21) Application No 9811466.3

(22) Date of Filing 28.05.1998

(71) Applicant(s)

3Com Technologies
(Incorporated in the Cayman Islands)
PO Box 309, Upland House, Georgetown,
Grand Cayman, Cayman Islands

(72) Inventor(s)

Tadhg Creedon
Paul James Moran
Graeme James McKerrel

(74) Agent and/or Address for Service

Bowles Horton
Felden House, Dower Mews, High Street,
BERKHAMSTED, Herts, HP4 2BL, United Kingdom

(51) INT CL⁶

H04L 12/56, H04Q 11/04

(52) UK CL (Edition Q)

H4P PPS

(56) Documents Cited

EP 0766425 A2 WO 97/43869 A1

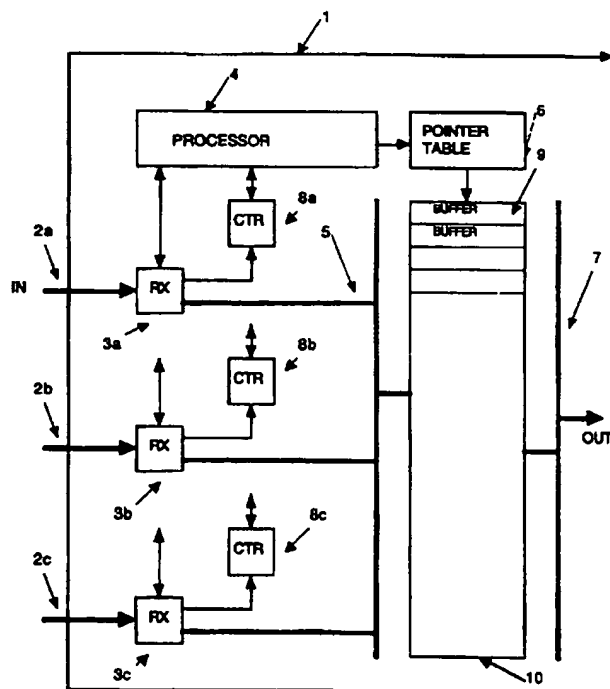
(58) Field of Search

UK CL (Edition P) H4K KTK, H4P PFD PPS POA PT
INT CL⁶ H04L 12/56
Online: WPI

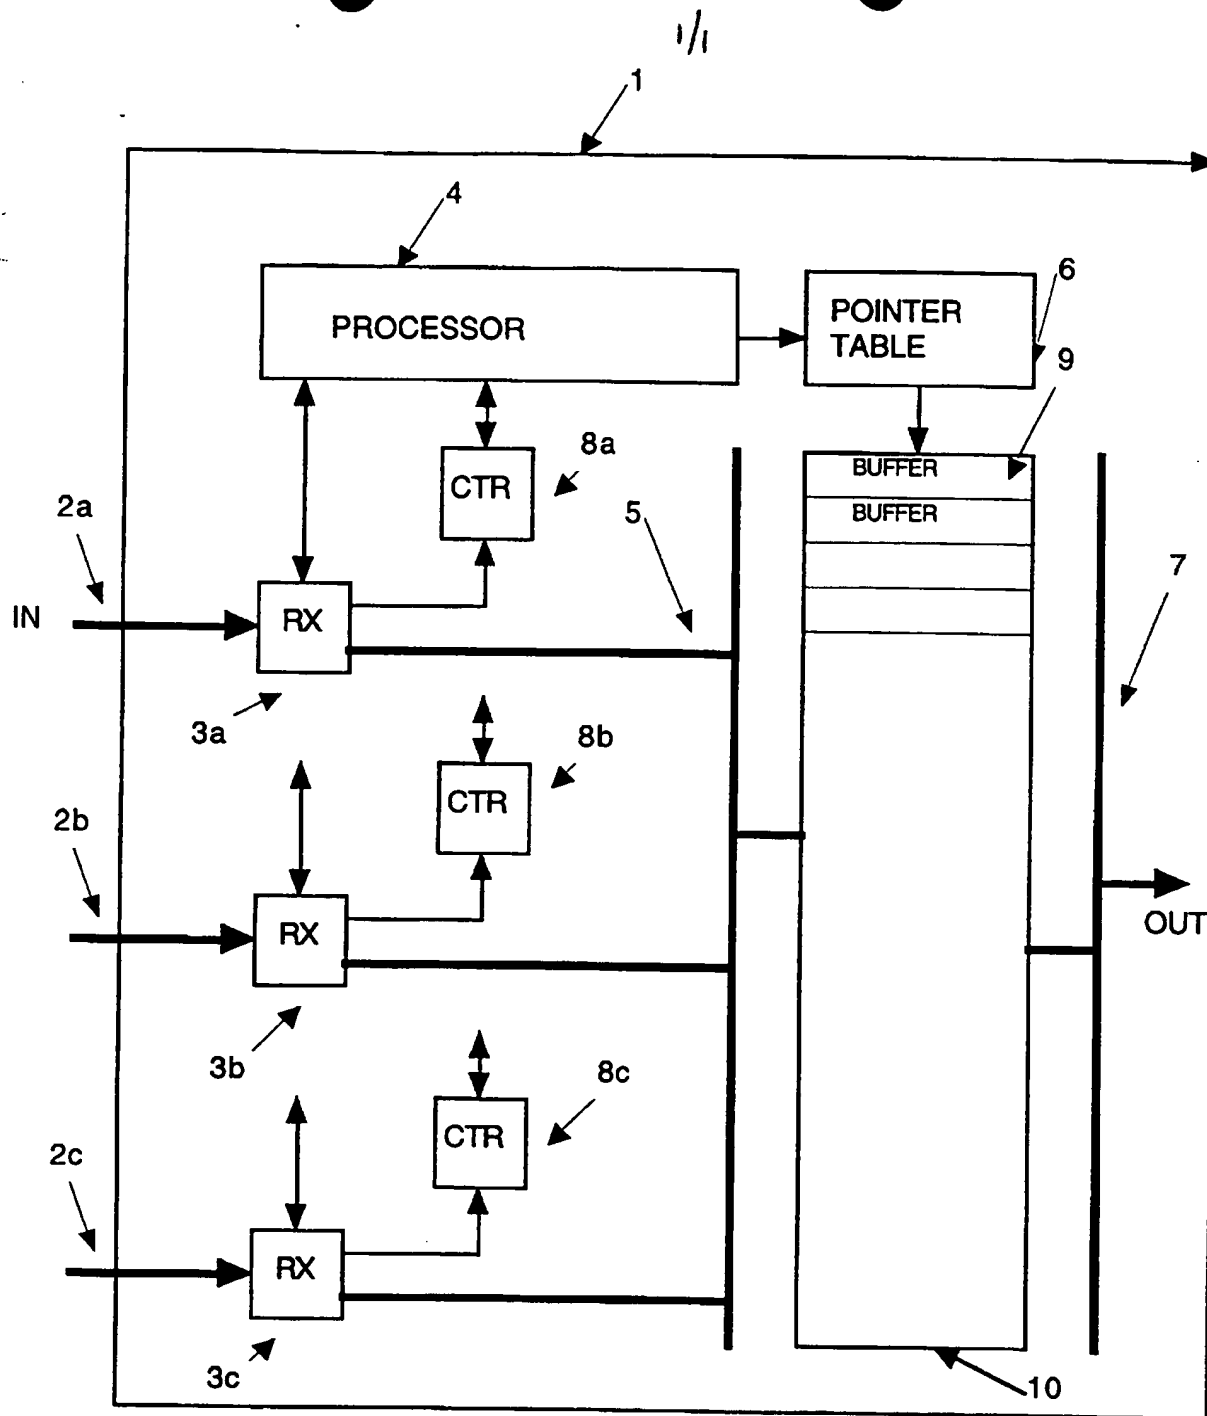
(54) Abstract Title

Buffer management in network devices

(57) A network device for use in a data communication network comprises: a multiplicity of ports (2a, 2b, 2c) capable of receiving input data; a common buffer memory (10) for storing in respective buffers (9) data received by the input ports; and processing means (4) for controlling the writing of data into buffers and the reading of data from said buffers. The processing means determines whether the extent of occupation of the common buffer memory exceeds a threshold defined to represent congestion of the memory. Counter means (8a, 8b, 8c) enabled by the processing means on detection of the aforementioned threshold count units of input data through the respective ports subsequent to the congestion threshold, the processing means decrementing these counters as data is read out from the common buffer memory. The counting levels may determine flow control or memory allocation.



GB 2 337 905 A



BUFFER MANAGEMENT IN NETWORK DEVICES

FIELD OF THE INVENTION

5 The present invention relates to the management of data storage in a buffer which forms part of a network device, such as a switch or repeater, which has a multiplicity of ports for the reception of data, usually data in the form of addressed packets. The invention particularly relates to an improved system of management intended to be useful where one or more of the ports that receive data is constantly busy and might utilise too much of the available memory
10 space in the device.

BACKGROUND TO THE INVENTION

15 Network devices which have a multiplicity of input ports for the reception of data such as packet data, from a variety of sources commonly provide buffering of the input ports. Thus data received at a port is stored in a buffer memory temporarily until such time as it can be transmitted onwards to its intended destination. By way of example only, the device may include some system for giving data of one type or another or data from a source or group of sources priority over other data. For this purpose the network device might have a contention-
20 resolution system, such as a csma/cd contention-resolution system wherein data from only one port at a time may be transmitted onwardly.

25 Modern network devices of this general character commonly include a single large memory, or possibly a set of memories which are operated as a single logical entity. The memory may be constituted by random access memory which is under the control of an appropriate processing device which may have recourse to a table of pointers which define and provide access to individual buffers within the memory. Data from any given port may be directed to a buffer by means of a respective pointer, this buffer remaining occupied until, normally under control of the processing means and a respective read pointer, the data is onwardly transmitted
30 either directly or by way of an output buffer and the buffer is "returned", that is to say made available for the storage of fresh data received from one or other of the input ports.

In systems of this nature, there is a significant danger that ports that are constantly busy can demand data storage in all the available memory space so that other ports which are receiving

data are deprived of buffers for the storage of data received at those ports. In such circumstances data packets may be dropped, to the obvious detriment of the performance of the network, particularly in respect of traffic through non-congested ports.

5 It is common practice to provide some protection for the congestion of buffer memory. One example is to provide for each receiving port a respective buffer of which the operation is independent of the main common buffer memory and which acts as another stage of buffering. For example, a first-in, first-out (FIFO) memory may be provided for each port. Such a solution requires additional hardware and control for each port. Another solution is to
10 associate in the main common memory certain groups of buffers with each port, so that a portion of the buffer memory is reserved for the use of each respective port. Such a system is more complex to control than an ordinary buffer memory and is somewhat inflexible.

SUMMARY OF THE INVENTION

15 In a buffer management system according to the invention, each port is associated with a counter, which is preferably normally cleared. If the common buffer memory approaches a congested state, which may be monitored by a processing means allocating space in the memory, these counters may be enabled to count. As buffers are allocated for data which
20 continues to arrive on the ports, the counters may be incremented, counting up or down according to the organisation of the counters, until certain watermarks represented by predetermined counts are reached. At the various watermarks, various actions can be taken, such as enabling flow control. The buffers may retain an identification of the receive port from which the data came and as data is transmitted, freeing the buffers, the respective counters may
25 be decremented. Buffers which are returned, that is to say made available for the reception of fresh received data, may be returned to a common pool or be allocated to the ports in a round robin fashion or according to some system of priority. If the memory becomes decongested at any stage, all the counters may be cleared.

BRIEF DESCRIPTION OF THE DRAWING

30 The drawing illustrates in schematic fashion a network switch according to the invention.

DETAILED DESCRIPTION

A network switch 1 according to the invention includes a multiplicity of receive ports of which only three (2a, 2b and 2c) are shown. In general there would be a substantial number, such as twenty-four or more ports for the network device. Each of the ports is connected through a receiving stage 3a - 3c respectively which is monitored and controlled by a processor 4. The ports are coupled by way of the receiving stages and a data bus 5 to a common buffer memory. In accordance with known practice, data packets or (more usually) groups of data packets are stored in buffers allocated to the data by means of "write" pointers under control of the processor. The pointers may be stored in a pointer table 6 and themselves control buffers 9 in the memory 10.

Data is read out from the buffers under the control of "read" pointers generated by the processor. Under what conditions the data is read out from the buffer memory is not of any great consequence to the invention and the common buffer memory for the input ports may only be a prelude to further stages within the same network device. For the purpose of the present invention it is only necessary that data stored in the buffer memory be read out from time to time under conditions which may vary according to the organisation or purpose of the device 1.

In the system as described, the buffer memory 10 is intended to provide a common pool of buffers for use by any of the ports. It follows that if one or more of the ports become unduly busy, it may devour too much of the available memory space so that other received ports are starved of buffers.

Associated with each of the ports is a respective counter (8a, 8b etc) which may be cleared by the processor 4 when the processor, which typically maintains a table of control pointers, detects when the buffer memory has reached some predetermined, selectable proportion of the total useable memory space. As buffers are allocated for data which continues to arrive on any given port, the respective counter is incremented. Each of the counters therefore indicates traffic in excess of a level which is set by an arbitrary definition of congestion or shortage of buffers within the buffer memory. The counters therefore provide an additional degree of control and, by setting various watermarks, i.e. particular counts which give rise to control signals if the count passes through those thresholds or watermarks, various actions, at the

choice of the designer, may be enabled. The detection of the predetermined counts, which may be programmable, and the consequential actions will be predetermined by the processor 4.

5 For example, a particular threshold may determine flow control so that if the traffic count for a particular port when the buffer memory is by definition likely to be congested exceeds some high watermark, the reception of data packets on that port may be inhibited. Alternatively, a small amount of memory space may be allocated to any port in which the traffic count after the onset of congestion in the buffer memory exceeds a watermark which is lower than the aforementioned high watermark.

10 Since the data in the buffer memories retains an identification of the port number, there is an easy means of enabling the processing means to decrement the traffic counter associated with a port when a buffer holding data previously received by this port is returned. However, the counters may be decremented according to a round robin scheme or other algorithm.

15 Various developments of the scheme are possible. For example, buffers may be allocated preferentially to ports which have a high traffic count, either on a predetermined or adaptive basis.

20

25

30

CLAIMS

1. A network device for use in a data communication network, and comprising:

5 a multiplicity of ports capable of receiving input data;

a common buffer memory for storing in respective buffers data received by the input ports;

10 processing means for controlling the writing of data into buffers and the reading of data from said buffers;

said processing means determining whether the extent of occupation of the common buffer memory exceeds a threshold defined to represent congestion of the memory; and

15 counter means enabled by the processing means on detection of the aforementioned threshold to count units of input data through the respective ports, the processing means decrementing these counters as data is read out from the common buffer memory.

20 2. A method of operating a network device which has a multiplicity of input ports and a common buffer memory for the temporary storage of data received by the input ports, any of the data ports being able to supply input data to the memory, the method comprising:

25 monitoring the occupation of the common memory to determine the exceeding of a threshold which indicates congestion in the memory;

enabling in response to said indication a respective counter associated with each input port so as to obtain a numerical representation of input data traffic through the respective port subsequent to the onset of congestion;

30 decrementing a counter when data has been read out from a buffer.

3. A method according to claim 2 wherein thresholds in the counting levels are employed to determine the further reception of input data at the respective input port or to control the allocation of available buffers for the storage of data received by the input ports.

5

10

15

20

25

30



Application No: GB 9811466.3
Claims searched: 1-3

Examiner: Keith Williams
Date of search: 16 November 1998

Patents Act 1977
Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK CI (Ed.P): H4P (PFD, PPS, PQA, PT); H4K (KTK)

Int CI (Ed.6): H04L 12/56

Other: Online WPI

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
A	EP 0766425 A2 Fujitsu Limited - see Fig.2; abstract	1,2
A	WO 97/43869 A1 Cisco Systems Inc. - see abstract	1,2

X Document indicating lack of novelty or inventive step
Y Document indicating lack of inventive step if combined with one or more other documents of same category.

& Member of the same patent family

A Document indicating technological background and/or state of the art.
P Document published on or after the declared priority date but before the filing date of this invention.
E Patent document published on or after, but with priority date earlier than, the filing date of this application.